arXiv:1612.02024v2 [math.ST] 9 Dec 2016

Impossible Inference in Econometrics: Theory and Applications to Regression Discontinuity, Bunching, and Exogeneity Tests

Marinho Bertanha *
University of Notre Dame

Marcelo J. Moreira
FGV

This version: December 12, 2016

First version: October 11, 2016

Abstract

This paper presents necessary and sufficient conditions for tests to have trivial power. By inverting these impractical tests, we demonstrate that the bounded confidence regions have error probability equal to one. This theoretical framework establishes a connection among many existing impossibility results in econometrics, those using the total variation metric and those using the Lévy-Prokhorov metric (convergence in distribution). In particular, the theory establishes conditions under which the two types of impossibility exist in econometric models. We then apply our theory to Regression Discontinuity Design models and exogeneity tests based on bunching.

Keywords: Hypothesis Tests, Confidence Intervals, Weak Identification, Regression Discontinuity

JEL Classification: C12, C14, C31

*Bertanha gratefully acknowledges support from CORE-UCL, and Moreira acknowledges the research support of CNPq and FAPERJ.

1 Introduction

The goal of most empirical studies is to estimate parameters of a population statistical model using a random sample of data. The difference between estimates and population parameters is uncertain because sample data do not have all the information about the population. Statistical inference provides methods for quantifying this uncertainty. Typical approaches include hypothesis testing and confidence sets. In a hypothesis test, the researcher divides all possible population models into two sets of models. The null set includes the models in which the researcher suspects may be false. The alternative set includes all other likely models. It is desirable to control the size of the test, that is, the error probability of rejecting the null set when the null set contains the true model. A powerful test has a small error probability of not rejecting the null set when the true model is outside the null set. Another approach is to use the data to build a confidence set for the unknown value of parameters of the true model. The researcher needs to control the error probability that the confidence set excludes the true value of parameters. This paper shows that the current implementation of hypothesis tests and confidence sets fails to control error probabilities in an important class of applications in economics.

Previous work demonstrates the impossibility of controlling error probabilities of tests and confidence sets in specific settings. There are essentially two types of impossibility results found in the literature. The first type of impossibility result says that any hypothesis test has power limited by size. That is, it is impossible to find a powerful test among tests with small error probability of mistakenly rejecting the null set. The second type of impossibility result states that any bounded confidence set has error probability equal to one. In words, it is impossible to find informative bounds on the true value of parameters that exclude the true value only with small probability.

In a classic paper, Bahadur and Savage (1956) show both types of impossibilities in the population mean case. Any test for distinguishing zero mean from nonzero mean distributions has power limited by size; and any bounded confidence interval for the population mean always has error probability equal to one. The key point is that any distribution with a nonzero mean is well-approximated by distributions with a zero mean. Bahadur and Savage (1956) employ the Total Variation (TV) metric to measure the distance between any two distributions. Romano (2004) makes an

important generalization of Bahadur and Savage (1956). Targeting hypothesis tests, Romano (2004) extends the first type of impossibility to the problem of testing any two sets of distributions that are indistinguishable in the TV metric. We refer to this notion of distance as *strong distance*.

A different branch of the econometrics literature focuses on the second type of impossibility of confidence sets. Gleser and Hwang (1987) consider classes of models indexed by finite dimensional parameters. They obtain the second type of impossibility whenever two sets of distributions are indistinguishable in the TV metric. Dufour (1997) generalizes the second type of impossibility of Gleser and Hwang (1987) to classes of models indexed by infinite dimensional parameters. Dufour (1997) also notes that the second type of impossibility implies that tests constructed from bounded confidence sets fail to control size. Unlike Bahadur and Savage (1956), and Romano (2004), Dufour (1997) relies on a notion of distance much weaker than the TV metric. A set of models $\mathbb{A}$ is indistinguishable from a set of models $\mathbb{B}$ if, for any model in set $\mathbb{A}$, there exists a sequence of models in set $\mathbb{B}$ that converges in distribution to that model in set $\mathbb{A}$. We refer to this notion of distance as *weak distance*.[1]

The main goal of this paper is to unite these two branches of the literature and show both types of impossibility results using the weak distance of Dufour (1997). We start with the problem of testing any two sets of distributions that are indistinguishable in the weak distance. We find that any test function almost-surely (a.s.) continuous in the data has power limited by size, the first type of impossibility. Next, we map distributions onto parameters of interest, and build confidence sets using a.s. continuous tests for these parameters. We demonstrate that bounded confidence sets have error probability equal to one, the second type of impossibility. As Dufour (1997) notes, tests that lead to bounded confidence sets also fail to control size. On the one hand, the weak distance does not allow us to limit power of every test function as Romano (2004) does with the TV metric. On the other hand, the class of a.s. continuous tests considered in this paper includes the vast majority of tests used in empirical studies. In that sense, the use of the TV metric is unnecessarily strong. It is

[1]Gleser and Hwang (1987) restrict their analysis to distributions that have parametric density functions with respect to the same sigma-finite measure. Two distributions are indistinguishable if their density functions are approximately the same pointwise in the data. In their setting, pointwise approximation in density functions is the same as approximation in the TV metric. However, pointwise approximation in density functions is still stronger than approximation in the weak distance of Dufour (1997). See Proposition 2.29 and Corollary 2.30, Van der Vaart (2000).

not a necessary condition for the second type of impossibility of testing either. For example, suppose we wish to test the null hypothesis that the distribution of a variable is discrete, versus the alternative that it is continuous. It is impossible to test these two sets of distributions. However, continuous distributions are not approximated by discrete distributions in the TV metric, as they are in the weak distance. We show that testing impossibility implies that the null and alternative sets of distributions are indistinguishable in the weak distance.

We then apply our reasoning to show that many empirical studies of models with discontinuities fail to control error probabilities. Numerous economic analyses identify parameters of interest by relying on natural discontinuities in the distribution of variables. This is the case of Regression Discontinuity Designs (RDD), an extremely popular identification strategy in economics. In RDD, the assignment of individuals into a program changes discontinuously at a cut-off point in a variable like age or test score, as in Hahn, Todd, and Van der Klaauw (2001) and Imbens and Lemieux (2008). For example, Schmieder, von Wachter, and Bender (2012) study individuals whose duration of unemployment insurance jumps with respect to age. Jacob and Lefgren (2004) look at students whose participation in summer school changes discontinuously with respect to test scores. Assuming all other characteristics vary smoothly at the cut-off, the effect of the summer school on future performance is captured by a discontinuous change in average performance at the cut-off. A fundamental assumption for identification is the *ceteris paribus* effect that performance varies smoothly with test scores. Models with continuous effects are well-approximated by models with discontinuous effects. Kamat (2015) uses the TV metric to show that the current practice of tests in RDD suffers from both the first type of impossibility. We revisit his result using the weaker notion of convergence, and we show the second type of impossibility also holds in RDD. The reasoning behind this finding arises from standard proofs for Portmanteau's theorem, in which indicator functions are approximated by continuous functions. On a positive note, the two types of impossibility theorems vanish if we restrict the class of functions/models, as in Armstrong and Kolesar (2015) (e.g., continuous functions with bounded derivatives).

In other applications, researchers assume a discontinuous change in unobserved characteristics of individuals at given points. This is the idea of bunching, widely exploited in economics. Bunching may occur because of a discontinuous change in incentives or a natural restriction on variables. For example, the distribution of indi-

viduals with respect to reported income may change discontinuously at points where income tax rates change, as in Saez (2010); or, the distribution of mothers, conditional on number of cigarettes smoked, may change discontinuously at zero smoking. Caetano (2015) proposes an exogeneity test for smoking as a determinant of birth weight that relies on bunching as evidence of endogeneity. A crucial assumption is the *ceteris paribus* effect that birth weight varies smoothly with smoking. Under this assumption, bunching is equivalent to the observed average birth weight being discontinuous at zero smoking. The exogeneity test looks for such discontinuity as evidence of endogeneity. Our point is that models in which birth weight is highly sloped or even discontinuous based on smoking are indistinguishable from smooth models. Therefore, we find the exogeneity test has power limited by size. The current implementation of tests for the size of discontinuity leads to bounded confidence sets so it also fails to control size. This example, applying the two types of impossibility theorems, appears to be novel in the literature.

The rest of this paper is divided as follows. Section 2 sets up a statistical framework for testing and building confidence sets. It presents the impossibility of controlling error probabilities for models that are close in the weak distance. Section 3 finds both types of inference impossibilities in economic models with discontinuities. Section 4 concludes.

2 Impossible Inference

The researcher has a sample of n iid observations $Z_i \sim P$, $Z_i \in \mathbb{R}^l$, $i = 1, \ldots, n$. The set of all possible probability distributions is $\mathbb{P}$. Every probability distribution $P \in \mathbb{P}$ is defined on the same sample space $\mathbb{R}^l$ and the Borel sigma-algebra $\mathcal{B}$. We are interested in testing the null hypothesis $H_0 : P \in \mathbb{P}_0$ versus the alternative hypothesis $H_1 : P \in \mathbb{P}_1$ for a partition $\mathbb{P}_0, \mathbb{P}_1$ of $\mathbb{P}$. We characterize a hypothesis test by a function of the data $\phi : \mathbb{R}^l \times \ldots \times \mathbb{R}^l \to [0, 1]$. If ϕ takes on only the values 0 and 1, the test is said to be *nonrandomized*. Given a sample $Z = (Z_1, \ldots, Z_n)$, we reject the null H_0 if the function $\phi(Z)$ equals one, but we fail to reject H_0 if $\phi(Z) = 0$. If the function $\phi(Z)$ is between 0 and 1, we reject the null with probability $\phi(Z)$ conditional on Z. The unconditional probability of rejecting the null hypothesis under distribution $P \in \mathbb{P}$ is denoted $\mathbb{E}_P[\phi]$.

The size of the test ϕ is $\sup_{P \in \mathbb{P}_0} \mathbb{E}_P[\phi]$. The power of the test under distribu-

tion $Q \in \mathbb{P}_1$ is given by $\mathbb{E}_Q[\phi]$. We say a test ϕ has power limited by size when $\sup_{Q\in\mathbb{P}_1} \mathbb{E}_Q[\phi] \leq \sup_{P\in\mathbb{P}_0} \mathbb{E}_P[\phi]$. We say an alternative distribution $Q \in \mathbb{P}_1$ is *powerfully detectable* if for every sequence of null distributions $\{P_k\}_{k=1}^{\infty}$ there exists a test ϕ such that $\mathbb{E}_{P_k}[\phi] \not\to \mathbb{E}_Q[\phi]$. There are two kinds of testing impossibilities in this section: strong and weak. *Strong testing impossibility* means no alternative distribution Q is powerfully detectable. *Weak testing impossibility* says that it is impossible to find a test ϕ with power greater than size. The strong testing impossibility implies the weak testing impossibility.

Lemma 1. *Suppose $\mathbb{P}_0$ and $\mathbb{P}_1$ are such that the strong testing impossibility holds. That is, $\forall Q \in \mathbb{P}_1$, $\exists\{P_k\}_{k=1}^{\infty}$ in $\mathbb{P}_0$ such that for any test ϕ, $\mathbb{E}_{P_k}[\phi] \to \mathbb{E}_Q[\phi]$. Then, the weak testing impossibility holds. That is, for any test ϕ, $\sup_{Q\in\mathbb{P}_1} \mathbb{E}_Q[\phi] \leq \sup_{P\in\mathbb{P}_0} \mathbb{E}_P[\phi]$.*

Proof of Lemma 1. Suppose the strong testing impossibility holds. Pick an arbitrary $Q \in \mathbb{P}_1$. There exists a sequence of null distributions $\{P_k\}_{k=1}^{\infty}$ such that for any test ϕ, $\mathbb{E}_{P_k}[\phi] \to \mathbb{E}_Q[\phi]$. Fix an arbitrary test ϕ.

Take an arbitrary sequence $\varepsilon_n \to 0$, and pick a subsequence P_{k_n} from the sequence above such that

$$-\varepsilon_n \leq \mathbb{E}_Q\phi - \mathbb{E}_{P_{k_n}}\phi \leq \varepsilon_n. \tag{2.1}$$

Therefore,

$$\mathbb{E}_Q\phi \leq \mathbb{E}_{P_{k_n}}\phi + \varepsilon_n \leq \sup_{P\in\mathbb{P}_0} \mathbb{E}_P\phi + \varepsilon_n. \tag{2.2}$$

Given $\varepsilon_n \to 0$, it follows that, for $\forall Q \in \mathbb{P}$,

$$\mathbb{E}_Q\phi \leq \sup_{P\in\mathbb{P}_0} \mathbb{E}_P\phi. \tag{2.3}$$

Consequently,

$$\sup_{Q\in\mathbb{P}_1} \mathbb{E}_Q\phi \leq \sup_{P\in\mathbb{P}_0} \mathbb{E}_P\phi, \tag{2.4}$$

as we wanted to prove. □

The similarity between models in $\mathbb{P}_0$ and $\mathbb{P}_1$ determines testing impossibility. There exist various notions of distance to measure the difference between two distributions P and Q. A common choice in the literature on testing impossibility is the

Total Variation (TV) metric $d_{TV}(P, Q)$:

$$d_{TV}(P, Q) = \sup_{B \in \mathcal{B}} |P(B) - Q(B)| . \tag{2.5}$$

The set of distributions $\mathbb{P}_1$ is said to be *indistinguishable* from $\mathbb{P}_0$ in the TV metric if, for any $Q \in \mathbb{P}_1$, there exists a sequence $\{P_k\}_{k=1}^{\infty}$ in $\mathbb{P}_0$ such that $d_{TV}(P_k, Q) \to 0$. Romano (2004) finds that any test ϕ has power limited by size as long as $\mathbb{P}_1$ is indistinguishable from $\mathbb{P}_0$ in the Total Variation (TV) metric. This striking impossibility result applies uniformly over the class of all hypothesis tests of the form $\phi(Z)$. The strong testing impossibility also holds when $\mathbb{P}_1$ is indistinguishable from $\mathbb{P}_0$ in the TV metric.

A weaker notion of distance relies on convergence in distribution. We say a sequence $\{P_k\}_{k=1}^{\infty}$ converges in distribution to Q, if, for every $B \in \mathcal{B}$ such that $Q(\partial B) = 0$, $P_k(B) \to Q(B)$. Here, ∂B is the boundary of a Borel set B, that is, the closure of B minus the interior of B. We denote convergence in distribution by $P_k \xrightarrow{d} Q$. This is equivalent to convergence in the Lévy-Prokhorov metric. Convergence of P_k to Q in the TV metric implies convergence in distribution. The converse does not hold.[2] In some settings, it is impossible to test the difference between two sets of distributions that are distinguishable in the TV metric but indistinguishable in the weak distance. For example, consider the null hypothesis that the distribution of a variable is discrete, versus the alternative that it is continuous. Any hypothesis test has power limited by size. Nevertheless, continuous distributions are not approximated by discrete distributions in the TV metric as they are in the weak distance. Therefore, the use of the TV metric is unnecessarily strong. In an influential paper, Chamberlain (1987) proves the asymptotic efficiency of the method of moments estimator. The proof relies on the approximation of discrete distributions by continuous distributions. This approximation is feasible with the weak distance, but unfeasible with the TV metric.

In this paper, we follow Dufour (1997) and use convergence in distribution rather than convergence in the TV metric to characterize indistinguishable sets of distribu-

[2]For example, a standardized binomial variable converges in distribution to a standard normal as the number of trials goes to infinity and the probability of success is fixed. It does not converge in the TV metric because the distance between these two distributions is always equal to one. In fact, consider the event equal to the entire real line minus the support of the binomial distribution. This event has unit probability under the normal distribution, but zero probability under the binomial distribution.

tions.

Assumption 1. *$\mathbb{P}_1$ is indistinguishable from $\mathbb{P}_0$ in the weak distance. That is, for every $Q \in \mathbb{P}_1$, there exists a sequence $\{P_k\}_{k=1}^{\infty}$ in $\mathbb{P}_0$ such that $P_k \xrightarrow{d} Q$.*

Assumption 1 is a necessary and sufficient condition for strong testing impossibility, as described in Theorem 1.

Theorem 1. *Consider the class of hypothesis tests $\phi(Z)$ that are almost-surely (a.s.) continuous in Z under any $Q \in \mathbb{P}_1$. If Assumption 1 holds, then the strong and weak testing impossibilities hold. Conversely, if the strong testing impossibility holds, then Assumption 1 holds.*

Remark 1. *The class of tests that are almost-surely (a.s.) continuous in Z under any $Q \in \mathbb{P}_1$ can be very large. For example, take a test that rejects the null when a test statistic is larger than a critical value: $\phi(Z) = I(\psi(Z) > c)$. This test is almost-surely (a.s.) continuous if the function ψ is continuous and $Q \in \mathbb{P}_1$ is absolutely continuous with respect to the Lebesgue measure.*

Remark 2. *The necessity part of Theorem 1 is true for all tests of the form $\phi(Z)$.*

Remark 3. *We do not need to restrict Theorem 1 to the class of a.s. continuous tests in every case. For example, consider $\mathbb{P}$ to be a subset of the parametric exponential family of distributions with parameter θ of finite dimension. Then, for any test ϕ, the power function of ϕ is continuous in θ, and Theorem 1 applies under Assumption 1. (see Theorem 2.7.1, Lehmann and Romano (2005)).*

Remark 4. *In many instances, Assumption 1 holds in both directions. That is, $\mathbb{P}_1$ is indistinguishable from $\mathbb{P}_0$, and $\mathbb{P}_0$ is indistinguishable from $\mathbb{P}_1$ in the weak distance. For example, Bahadur and Savage (1956) find that any distribution with mean m is well-approximated by distributions with mean $m' \neq m$, and vice-versa. Section 3.1 finds the same bidirectionality for models with or without discontinuities. If Assumption 1 holds in both directions, switching the roles of $\mathbb{P}_0$ and $\mathbb{P}_1$ in Theorem 1 shows that power is equal to size.*

Proof of Theorem 1. The proof of the sufficiency part of Theorem 1 follows the same lines as the proof of Theorem 1 by Romano (2004).

Fix $Q \in \mathbb{P}_1$, and pick a sequence P_k in $\mathbb{P}_0$ such that $P_k \xrightarrow{d} Q$. Convergence in distribution $P_k \xrightarrow{d} Q$ is equivalent to $\mathbb{E}_{P_k}[g] \to \mathbb{E}_Q[g]$ for every bounded real-valued

function g whose set of discontinuity points has probability zero under Q (Theorem 25.8, Billingsley (2008)). In particular, this is true for $g = \phi$, for any ϕ that is a.s. continuous under Q. This shows the strong testing impossibility. The weak testing impossibility follows by Lemma 1.

For the converse, suppose the strong testing impossibility holds. Then, $\forall Q \in \mathbb{P}_1$, there exists a sequence P_k in $\mathbb{P}_0$ such that $\forall \phi$, $\mathbb{E}_{P_k}[\phi] \to \mathbb{E}_Q[\phi]$. The same is true for every bounded real-valued function g that is a.s. continuous under Q. Therefore, $P_k \xrightarrow{d} Q$, and Assumption 1 holds. □

It is useful to connect testing impossibility with the impossibility of controlling error probability of confidence sets found by Gleser and Hwang (1987) and Dufour (1997). Define a real-valued function $\mu : \mathbb{P} \to \mathbb{R}$, for example, mean, variance, median, etc. The set of distributions $\mathbb{P}$ is implicitly chosen such that μ is well-defined. We consider real-valued functions for simplicity, and results for μ with more general ranges are straightforward to obtain. The range of μ is $\mu(\mathbb{P})$. Suppose we are interested in a confidence set for $\mu(P)$ when the true model is $P \in \mathbb{P}$. A confidence set takes the form of a function $C(Z)$ of the data. For a model $P \in \mathbb{P}$, the *coverage probability* of $C(Z)$ is given by $\mathbb{P}_P[\mu(P) \in C(Z)]$. The confidence region $C(Z)$ has *confidence level* $1 - \alpha$ (i.e. error probability α) if $C(Z)$ contains $\mu(P)$ with probability at least $1 - \alpha$:

$$\inf_{P \in \mathbb{P}} \mathbb{P}_P[\mu(P) \in C(Z)] = 1 - \alpha \tag{2.6}$$

The function $C(Z)$ is constructed by inverting a test in the following manner. For a given $m \in \mu(\mathbb{P})$, define

$$\mathbb{P}_0(m) = \{P \in \mathbb{P} : \mu(P) = m\} \text{ and} \tag{2.7}$$

$$\mathbb{P}_1(m) = \mathbb{P} \cap \mathbb{P}_0(m)^c \tag{2.8}$$

where $\mathbb{P}_0(m)^c$ denotes the complement of the set $\mathbb{P}_0(m)$. Let $\phi_m(Z)$ be a nonrandomized test for testing $H_0 : P \in \mathbb{P}_0(m)$ against $H_1 : P \in \mathbb{P}_1(m)$. Then,

$$C(Z) = \{m \in \mu(\mathbb{P}) : \phi_m(Z) = 0\}. \tag{2.9}$$

Lemma 2 gives the relationship between the coverage probability of $C(Z)$ and the size of $\phi_m(Z)$.

Lemma 2. *Let $C(Z)$ be constructed as in Equation (2.9). The confidence level of $C(Z)$ is equal to one minus the supremum of the sizes of the tests $\phi_m(Z)$ over $m \in \mu(\mathbb{P})$.*

Proof of Lemma 2. Suppose

$$\sup_{m \in \mu(\mathbb{P})} \sup_{P \in \mathbb{P}_0(m)} \mathbb{P}_P(\phi_m(Z) = 1) = \alpha. \tag{2.10}$$

Now, pick $\varepsilon > 0$. Then, $\exists m_\varepsilon$ such that

$$\alpha - \varepsilon/2 \leq \sup_{P \in \mathbb{P}_0(m_\varepsilon)} \mathbb{P}_P(\phi_{m_\varepsilon}(Z) = 1) \leq \alpha. \tag{2.11}$$

There also $\exists P_\varepsilon \in \mathbb{P}_0(m_\varepsilon)$ such that

$$\alpha - \varepsilon \leq \mathbb{P}_{P_\varepsilon}(\phi_{m_\varepsilon}(Z) = 1) \leq \alpha. \tag{2.12}$$

Rearranging the expression above, we obtain

$$1 - \alpha + \varepsilon \geq \mathbb{P}_{P_\varepsilon}(\mu(P_\varepsilon) \in C(Z)) \geq 1 - \alpha. \tag{2.13}$$

Therefore, we find that

$$\inf_{P \in \mathbb{P}} \mathbb{P}_P \left[\mu(P) \in C(Z)\right] = 1 - \alpha, \tag{2.14}$$

as we wanted to proven. □

We apply Theorem 1 to every $m \in \mu(\mathbb{P})$ to show the impossibility of controlling coverage probabilities found by Gleser and Hwang (1987) and Dufour (1997).

Theorem 2. *For every $m \in \mu(\mathbb{P})$, suppose Assumption 1 holds for $\mathbb{P}_0(m), \mathbb{P}_1(m)$ and that $\phi_m(Z)$ is a.s. continuous in Z under every $Q \in \mathbb{P}_1(m)$. If the confidence set $C(Z)$ of Equation (2.9) has confidence level $1 - \alpha$, then*

$$\mathbb{P}_P \left[m \in C(Z)\right] \geq 1 - \alpha \quad \forall P \in \mathbb{P} \quad \forall m \in \mu(\mathbb{P}). \tag{2.15}$$

Moreover, if $1 - \alpha > 0$, and $\mu(\mathbb{P})$ is unbounded, then the confidence set $C(Z)$ is unbounded with strictly positive probability for every $P \in \mathbb{P}$.

Remark 5. *The contrapositive statement of the second part of Theorem 2 says the following. If $\mu(\mathbb{P})$ is unbounded, then any confidence set that is a.s. bounded under some distribution $P \in \mathbb{P}$ has zero confidence level. In addition, the test $\phi_m(Z)$ fails to control size. In fact, by Lemma 2, there exists $m \in \mu(\mathbb{P})$ for which the size of $\phi_m(Z)$ is arbitrarily close to one.*

Remark 6. *Moreira (2003) provides numerical evidence that Wald tests can have large null rejection probabilities for the null of no causal effect ($m = 0$) in the simultaneous equations model. To show that Wald tests have null rejection probabilities arbitrarily close to one, the hypothesized value m for the null would need to change as well. He also suggests replacing the critical value by a critical value function of the data. This critical value function depends on the hypothesized value m. Our theory shows that this critical value function is unbounded if we were to change m freely.*

Proof of Theorem 2. Fix $m \in \mu(\mathbb{P})$. Theorem 1 holds for $\phi_m(Z)$, $\mathbb{P}_0(m)$, $\mathbb{P}_1(m)$, so that

$$\sup_{Q \in \mathbb{P}_1(m)} \mathbb{E}_Q \phi_m(Z) \leq \sup_{P \in \mathbb{P}_0(m)} \mathbb{E}_P \phi_m(Z). \tag{2.16}$$

By assumption,

$$\mathbb{P}_P[m \in C(Z)] \geq 1 - \alpha \text{ for all } P \in \mathbb{P}_0(m). \tag{2.17}$$

This is equivalent to

$$\mathbb{P}_P[m \in C(Z)] = \mathbb{P}_P[\phi_m(Z) = 0] = 1 - \mathbb{E}_P[\phi_m(Z)] \geq 1 - \alpha. \tag{2.18}$$

In short,

$$\mathbb{E}_P[\phi_m(Z)] \leq \alpha \text{ for all } P \in \mathbb{P}_0(m). \tag{2.19}$$

Using Expression (2.16),

$$\mathbb{E}_Q \phi_m(Z) \leq \alpha \quad \forall Q \in \mathbb{P}_1(m). \tag{2.20}$$

Then, for any $P \in \mathbb{P}$,

$$\mathbb{E}_P[\phi_m(Z)] \leq \alpha. \tag{2.21}$$

This implies that

$$\mathbb{P}_P[\phi_m(Z) = 0] \geq 1 - \alpha \tag{2.22}$$

and, consequently,

$$\mathbb{P}_P[m \in C(Z)] \geq 1 - \alpha, \tag{2.23}$$

which proves the first part of the theorem.

For the second part, if $\mu(\mathbb{P})$ is unbounded, then any arbitrarily large $m \in \mu(\mathbb{P})$ is covered by the confidence set $C(Z)$ with strictly positive probability under any $P \in \mathbb{P}$. $\square$

3 Applications

In this section, we give three important examples of economic applications of models with discontinuities. Inference with bounded confidence sets of correct coverage or tests of power larger than size is impossible in these applications.

3.1 Regression Discontinuity Designs

The first example is the Regression Discontinuity Design (RDD), first formalized by Hahn, Todd, and Van der Klaauw (2001) (HTV). RDD has had an enormous impact in applied research in various fields of economics. Applications of RDD started gaining popularity in economics in the 1990s. Influential papers include Black (1999), who studies the effect of quality of school districts on house prices, where quality changes discontinuously across district boundaries; Angrist and Lavy (1999), who measure the effect of class sizes on academic performance, where size varies discontinuously with enrollment; and Lee (2008), who analyzes U.S. House elections and incumbency where election victory is discontinuous on the share of votes.

Recent theoretical contributions include the study of rate optimality of RDD estimators by Porter (2003) and the data-driven optimal bandwidth rules by Imbens and Kalyanaraman (2012) and Calonico, Cattaneo, and Titiunik (2014). RDD identifies causal effects local to a cut-off value and several authors developed conditions for extrapolating local effects away from the cut-off. These include estimation of derivatives of the treatment effect at the cut-off by Dong (2016) and Dong and Lewbel (2015); tests for homogeneity of treatment effects in Fuzzy RDD by Bertanha and Imbens (2016); and estimation of average treatment effects in RDD with variation in cut-off values by Bertanha (2016). All these theoretical contributions rely on point identification and inference, and they are subject to both types of impossibility results. The

current practice of testing and building confidence intervals relies on Wald test statistics $(t(Z) - m)/s(Z)$, where $t(Z)$ and $s(Z)$ are a.s. continuous and bounded in the data. For a choice of critical value z, hypothesis tests $\phi(Z) = \mathbb{I}\{|(t(Z)-m)/s(Z)| > z\}$ are a.s. continuous when the data is continuously distributed. Confidence intervals $C(Z) = \{t(Z) - s(Z)z \leq m \leq t(Z) + s(Z)z\}$ have a.s. bounded length $2s(Z)z$.

The usual setup of RDD follows the widely used Rubin Causal Model or Potential Outcome framework. For each individual $i = 1, \ldots, n$, define four primitive random variables $D_i, X_i, Y_i(0), Y_i(1)$. These variables are independent and identically distributed. The variable D_i takes values in $\{0, 1\}$ and indicates treatment status. The real-valued variables $Y_i(0)$ and $Y_i(1)$ denote the potential outcomes, respectively, if untreated and treated. Finally, the forcing variable X_i represents a real-valued characteristic of the individual that is not affected by the treatment. The forcing variable has a continuous probability density function $f(x)$ with interval support equal to $\mathbb{X}$. The econometrician observes X_i, D_i, but only one of the two potential outcomes for each individual: $Y_i = D_i Y_i(1) + (1 - D_i)Y_i(0)$. For simplicity, we consider the sharp RDD case, but it is straightforward to generalize our results to the fuzzy case. In the sharp case, agents receive the treatment if, and only if, the forcing variable is greater than or equal to a fixed policy cut-off c in the interior of support $\mathbb{X}$. Hence, $D_i = \mathbb{I}\{X_i \geq c\}$, where $\mathbb{I}\{\cdot\}$ denotes the indicator function.

We focus on average treatment effects. In RDD settings, identification of average effects is typically obtained only at the cut-off value after assuming continuity of average potential outcomes conditional on the forcing variable. In other words, we assume that $\mathbb{E}[Y_i(0)|X_i = x]$ and $\mathbb{E}[Y_i(1)|X_i = x]$ are bounded continuous functions of x. HTV show that this leads to identification of the parameter of interest:

$$m = \mathbb{E}\left[Y_i(1) - Y_i(0)|X_i = c\right] = \lim_{x \downarrow c} \mathbb{E}\left[Y_i|X_i = x\right] - \lim_{x \uparrow c} \mathbb{E}\left[Y_i|X_i = x\right]. \tag{3.1}$$

Let $\mathcal{G}$ denote the space of all functions $g : \mathbb{X} \to \mathbb{R}$ that are bounded, and that are continuous in every $x \in \mathbb{X} \setminus \{c\}$. The notation $\mathbb{X} \setminus \{c\}$ represents the set with every point of $\mathbb{X}$ except for c. Each individual pair of variables (Y_i, X_i) is i.i.d. as P. The family of all possible models for P is denoted as

$$\mathbb{P} = \{P : (Y_i, X_i) \sim P,\ \exists g \in \mathcal{G} \text{ s.t. } \mathbb{E}[Y_i|X_i = x] = g(x)\}. \tag{3.2}$$

The local average causal effect is the function of the distribution of the data $P \in \mathbb{P}$

Figure 1: Continuous Approximation of Discontinuous Conditional Mean

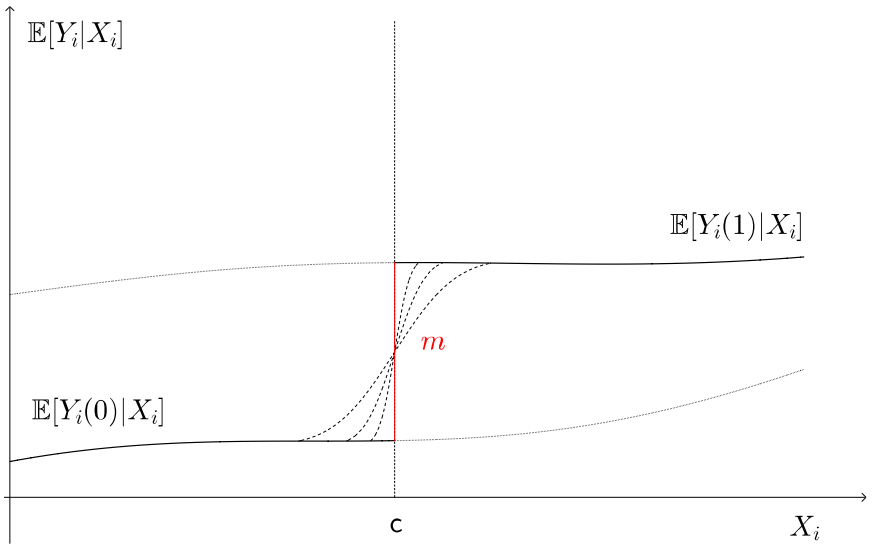

Notes: The graph depicts a positive treatment effect $m = \mathbb{E}[Y_i(1)|X_i] - \mathbb{E}[Y_i(0)|X_i]$. Dotted lines denote unobserved parts of $\mathbb{E}[Y_i(1)|X_i]$ and $\mathbb{E}[Y_i(0)|X_i]$. The solid line denotes $\mathbb{E}[Y_i|X_i]$, and has a jump discontinuity at $X_i = c$. The discontinuous conditional mean $\mathbb{E}[Y_i|X_i]$ is approximated by a sequence of continuous conditional mean functions (dashed lines) in the weak distance.

given by (3.1), provided the identification assumptions of HTV hold. The parameter m of the size of the discontinuity is weakly identified in the set of possible true models $\mathbb{P}$. Intuitively, any conditional mean (bounded) function $\mathbb{E}[Y_i|X_i = x]$ with a discontinuity at $x = c$ is well-approximated by a sequence of continuous conditional mean (bounded) functions. This is illustrated in Figure 1 below.

The reasoning behind Figure 1 easily verifies Assumption 1 for the RDD case.[3]

Corollary 1. *Assumption 1 is satisfied for $\forall m \in \mathbb{R}$, and Theorems 1 and 2 apply to the RDD case. Namely, (i) a.s. continuous tests $\phi_m(Z)$ on the value of the discontinuity m have power limited by size; (ii) confidence sets on the value of the discontinuity m and with finite expected length have zero confidence level.*

Remark 7. *Corollary 1 also applies to quantile treatment effects by simply changing the definition of the functional $\mu(\mathbb{P})$ to be the difference in side limits of a conditional*

[3]A similar proof technique is used to demonstrate part of Portmanteau's theorem. It is known that, if $\mathbb{E}[f(X_n)] \to \mathbb{E}[f(X)]$ for every bounded function f that is a.s. continuous under the distribution of X, then $X_n \xrightarrow{d} X$. The proof uses a continuous function f that is approximately equal to an indicator function, as in Figure 1. See Theorem 25.8, Billingsley (2008).

τ-th quantile $Q_\tau(Y_i|X_i = x)$ at $x = c$.

Remark 8. *In the fuzzy RDD case, the treatment effect is equal to the discontinuity in $\mathbb{E}[Y_i|X_i]$ at $X_i = c$ divided by the discontinuity in $\mathbb{E}[D_i|X_i]$ at $X_i = c$. Corollary 1 applies to both of these conditional mean functions, and it leads to impossible inference in the fuzzy RDD case as well. Feir, Lemieux, and Marmer (2016) study weak identification in Fuzzy RDD and propose a robust testing procedure. Differently from Kamat (2015) and this paper, their source of weak identification comes from an arbitrarily small discontinuity in $\mathbb{E}[D_i|X_i]$ at $X_i = c$.*

Proof of Corollary 1. Fix $m \in \mathbb{R}$. Pick an arbitrary $Q \in \mathbb{P}_1(m)$, and let $m' = \mu(Q) \neq m$. Construct a sequence of functions $g_k : \mathbb{R} \to \mathbb{R}$, $g_k \in \mathcal{G}$, $k = 1, 2, \ldots$ as follows:

$$g_k(x) = \mathbb{E}_Q[Y_i|X_i = x] + \mathbb{I}\{c - 1/k \leq x < c\}k(m' - m)(x - c + 1/k). \tag{3.3}$$

The function g_k is equal to the conditional mean function $\mathbb{E}_Q[Y_i|X_i = x]$ except for $x \in [c-1/k, c)$. Note also that $\lim_{x\downarrow c} g_k(x) - \lim_{x\uparrow c} g_k(x) = m$. Pick a sequence $P_k \in \mathbb{P}$ such that

1. the marginal distribution of X_i under P_k is the same as under Q;
2. the conditional distribution of $(Y_i, X_i)|X_i \notin [c-1/k, c)$ under P_k is the same as under Q; and
3. $\mathbb{E}_{P_k}[Y_i|X_i = x] = g_k(x)$.

Note that $\mu(P_k) = m$ and $P_k \in \mathbb{P}_0(m) \ \ \forall k$. All we have to show next is that $P_k \xrightarrow{d} Q$.

Define $A_k = \{c - k^{-1} \leq X_i < c\}$, and let A_k^c be the complement of A_k. Then, for any measurable event $\mathcal{A}$ of (Y_i, X_i),

$$\mathbb{P}_{P_k}[(X_i, Y_i) \in \mathcal{A}] = \mathbb{P}_{P_k}[(X_i, Y_i) \in \mathcal{A} \cap A_k] + \mathbb{P}_{P_k}[(X_i, Y_i) \in \mathcal{A} \cap A_k^c].$$

For the first term, we show that

$$0 \leq \mathbb{P}_{P_k}[(X_i, Y_i) \in \mathcal{A} \cap A_k] \leq \mathbb{P}_{P_k}[X_i \in A_k] = \mathbb{P}_Q[X_i \in A_k] = 0.$$

For the second term, we find that

$$\mathbb{P}_{P_k}[(X_i, Y_i) \in \mathcal{A} \cap A_k^c] = \mathbb{P}_{P_k}[(X_i, Y_i) \in \mathcal{A} \mid X_i \in A_k^c].\mathbb{P}_{P_k}[X_i \in A_k^c] \to \mathbb{P}_Q[(X_i, Y_i) \in \mathcal{A}],$$

by the continuity property of probability measures because $A_k \downarrow \emptyset$, where $\emptyset$ denotes the empty set, and $A_k^c \uparrow \mathbb{R}$. Therefore, Assumption 1 is satisfied for every $m \in \mathbb{R}$. Theorem 1 applies, and Theorem 2 applies with $\mu(\mathbb{P}) = \mathbb{R}$. □

The most common inference procedures currently in use in applied research with RDD rely on Wald tests that are a.s. continuous in the data and produce confidence intervals of finite expected length. See Imbens and Kalyanaraman (2012) and Calonico, Cattaneo, and Titiunik (2014) for most used inference procedures. Corollary 1 implies it is impossible to control size of these tests and coverage of these confidence intervals.

Ours is not the first paper to show impossible inference in the RDD case. Kamat (2015) demonstrates the important fact that models with a discontinuity are similar to models without a discontinuity in the TV metric. He applies the testing impossibility of Romano (2004) and finds that tests have power limited by size. Using the graphical intuition of Figure 1, we provide a simpler proof of the same facts, using the weak distance instead of the TV metric. Moreover, we add the result that confidence intervals produced from Wald tests have zero confidence level.

A second example of an application of Theorems 1 and 2 to a model with discontinuity is the so-called Regression Kink Design (RKD). RKD has recently gained popularity in economics. See, for example, Dong (2016), Nielsen, Sørensen, and Taber (2010), and Simonsen, Skipper, and Skipper (2016). The setup is the same as in the RDD case, except that the causal effect of interest is the change in the slope of the conditional mean of outcomes at the threshold. Continuity of the first derivatives $\nabla_x \mathbb{E}[Y_i(1)|X_i = x]$ and $\nabla_x \mathbb{E}[Y_i(0)|X_i = x]$ at the threshold $x = c$ guarantees identification of the average effect. The parameter of interest $m = \mu(P)$ is a function of the distribution of $Z_i = (Y_i, X_i)$:

$$\mu(P) = \nabla_x \mathbb{E}[Y_i(1) - Y_i(0)|X_i = x] = \lim_{x \downarrow c} \nabla_x \mathbb{E}[Y_i|X_i = x] - \lim_{x \uparrow c} \nabla_x \mathbb{E}[Y_i|X_i = x]. \quad (3.4)$$

The family of all possible distributions of Z_i is defined in a slightly different way

than in Equation (3.2):

$$\mathbb{P} = \{P : (Y_i, X_i) \sim P,\ \exists g \in \mathcal{G} \text{ s.t. } \nabla_x \mathbb{E}[Y_i|X_i = x] = g(x)\}. \tag{3.5}$$

Weak identification of μ arises from the fact that any conditional mean function $\mathbb{E}[Y_i|X_i = x]$ with a discontinuous first derivative at $x = c$ is well-approximated by a sequence of continuously differentiable conditional mean functions. Assumption 1 is easily verified using this insight.

Corollary 2. *Assumption 1 is satisfied for $\forall m \in \mathbb{R}$, and Theorems 1 and 2 apply to RDD. Namely, (i) a.s. continuous tests $\phi_m(Z)$ on the value of the kink discontinuity m have power limited by size; (ii) confidence sets on the value of the kink discontinuity m and with finite expected length have zero confidence level.*

Proof of Corollary 2. Our proof follows that of Corollary 1. Simply use the new definitions of $\mathbb{P}$ and $\mu(P)$. Construct the sequence P_k with $\nabla_x \mathbb{E}_{P_k}[Y_i|X_i = x] = g_k(x)$. □

3.2 Exogeneity Tests Based on Discontinuities

A third example of an application of Theorems 1 and 2 to a model with discontinuity comes from the idea of bunching, widely exploited in economics. Bunching occurs when the distribution of individual types conditional on a scalar variable X changes discontinuously at some known value of X. For example, the US income tax schedule has kinks at certain income values. The distribution of individuals conditional on income may change discontinuously with respect to income (Saez (2010)).

Caetano (2015) uses the idea of bunching to construct an exogeneity test that does not require instrumental variables. It applies to regression models where unobserved endogenous factors are assumed to be discontinuous with respect to an explanatory variable. Of interest is the impact of a scalar explanatory variable X on an outcome variable Y after controlling for covariates W. The variable X is potentially endogenous. For example, suppose we are interested in the effect of average number of cigarettes smoked per day X on birth weight Y, after controlling for mothers's observed characteristics W.

The population model that determines Y is written as $Y = h(X, W) + U$, where U summarizes unobserved confounding factors affecting Y. We are unable to observe

bunching unless h is assumed continuous on (X, W). Bunching of U with respect to X is evidence of endogeneity of X. In the birth weight example, researchers typically assume the distribution of mothers's unobserved characteristics U changes discontinuously from the non-smoking $X = 0$ to the smoking $X > 0$ populations. This is the same as bunching at $X = 0$ or local endogeneity of X at 0. Bunching at 0 implies discontinuity of $\mathbb{E}[U|X = 0, W] - \mathbb{E}[U|X = x, W]$ at $x = 0$. Continuity of h makes bunching equivalent to a discontinuity of $\mathbb{E}[Y|X = 0, W = w] - \mathbb{E}[Y|X = x, W = w]$ at $x = 0$ for every w. Caetano (2015) proposes testing

$$\mathbb{E}[Y|X = 0, W = w] - \mathbb{E}[Y|X = x, W = w] = 0 \ \forall w \tag{3.6}$$

as a means of testing for local exogeneity of X at $X = 0$. We argue that h may have a high slope on X or even be discontinuous on X which makes exogeneity untestable.

The observed data $Z = (Z_1, \ldots, Z_n)$, $Z_i = (Y_i, X_i, W_i)$ is i.i.d. with probability P. The scalar variable X_i is continuously distributed, and the support of (X_i, W_i) is denoted $\mathbb{X} \times \mathbb{W}$. Assume $\exists \delta > 0$ such that $[0, \delta) \subset \mathbb{X}$. Let $\mathcal{G}$ denote the space of all functions $g : \mathbb{X} \times \mathbb{W} \to \mathbb{R}$ that are bounded, and that are continuous in every $x \in \mathbb{X} \setminus \{0\}$. The family of all possible distributions P is denoted as

$$\mathbb{P} = \{P : Z_i \sim P, \ \exists g \in \mathcal{G} \text{ s.t. } \mathbb{E}_P[Y_i|X_i = x, W_i = w] = g(x, w)\}. \tag{3.7}$$

Under local exogeneity of X, the function $\tau_P(w) = \mathbb{E}_P[Y_i|X_i = 0, W_i = w] - \lim_{x \downarrow 0} \mathbb{E}[Y_i|X_i = x, W_i = w]$ must be equal to 0 $\forall w \in \mathbb{W}$. In practice, it is convenient to conduct inference on an aggregate of $\tau_P(w)$ over $w \in \mathbb{W}$ instead of on the entire function $\tau_P(w)$. Examples of aggregation include the average of $|\tau_P(w)|$, the square root of the average of $\tau_P(w)^2$, or the supremum of $|\tau_P(w)|$ over $\mathbb{W}$. For a fixed choice of aggregation, assume that if $\tau_P(w)$ is equal to a constant m $\forall w \in \mathbb{W}$, then the aggregated value of $\tau_P(w)$ over $w \in \mathbb{W}$ is also equal to that same constant m. For a distribution $P \in \mathbb{P}$, denote the aggregated value of $\tau_P(w)$ as $\mu(P)$. The local exogeneity corresponds to the test of $\mu(P) = 0$ versus $\mu(P) \neq 0$.

The parameter $\mu(P)$ is weakly identified in the class of models $\mathbb{P}$. Just as in the RDD case, any conditional mean function $\mathbb{E}[Y_i|X_i = x, W_i = w]$ with a discontinuity at $x = 0$ is well-approximated by a sequence of continuous conditional mean functions $\mathbb{E}[Y_i|X_i = x, W_i = w]$. Assumption 1 is verified using the same argument as the RDD case.

Corollary 3. *Assumption 1 is satisfied for $\forall m \in \mathbb{R}$, and Theorems 1 and 2 apply to the case of the local exogeneity test. Namely, (i) a.s. continuous tests $\phi_m(Z)$ on the value of the aggregate discontinuity m have power limited by size; (ii) confidence sets on the value of the aggregate discontinuity m and with finite expected length have zero confidence level.*

Proof of Corollary 3. Fix $m \in \mu(\mathbb{P})$. Pick an arbitrary $Q \in \mathbb{P}_1(m)$, and let $m' = \mu(Q) \neq m$. Construct a sequence of functions $g_k : \mathbb{X} \times \mathbb{W} \to \mathbb{R}$, $g_k \in \mathcal{G}$, $k = 1, 2, \ldots$ as follows:

$$g_k(x,w) = \mathbb{E}_Q[Y_i|X_i = x, W_i = w] + \mathbb{I}\{0 < x < 1/k\}$$
$$\left[\mathbb{E}_Q[Y_i|X_i = 0, W_i = w] - \lim_{x \downarrow 0} \mathbb{E}_Q[Y_i|X_i = x, W_i = w] - m\right](1 - kx).$$

The function g_k is equal to the conditional mean function $\mathbb{E}_Q[Y_i|X_i = x, W_i = w]$ except for $x \in (0, 1/k)$. Note also that $g_k(0,w) - \lim_{x \downarrow 0} g_k(x,w) = m$. Pick a sequence $P_k \in \mathbb{P}$ such that:

1. the marginal distribution of X_i under P_k is the same as under Q;
2. the conditional distribution of $(Y_i, X_i, W_i)|X_i \notin (0, 1/k)$ under P_k is the same as under Q; and
3. $\mathbb{E}_{P_k}[Y_i|X_i = x, W_i = w] = g_k(x,w)$.

Note that $\mu(P_k) = m$, because $\tau_{P_k}(w) = m$ $\forall w$, and $P_k \in \mathbb{P}_0(m)$ $\forall k$. All we have to show next is that $P_k \xrightarrow{d} Q$.

Define $A_k = \{0 < X_i < 1/k\}$, and let A_k^c be the complement of A_k. Then, for any measurable event $\mathcal{A}$ of (Y_i, X_i, W_i),

$$\mathbb{P}_{P_k}[(Y_i, X_i, W_i) \in \mathcal{A}] = \mathbb{P}_{P_k}[(Y_i, X_i, W_i) \in \mathcal{A} \cap A_k] + \mathbb{P}_{P_k}[(Y_i, X_i, W_i) \in \mathcal{A} \cap A_k^c]. \quad (3.8)$$

For the first term, we show that

$$\mathbb{P}_{P_k}[(Y_i, X_i, W_i) \in \mathcal{A} \cap A_k] \leq \mathbb{P}_{P_k}[X_i \in A_k] = \mathbb{P}_Q[X_i \in A_k] = 0.$$

For the second term, we find that

$$\mathbb{P}_{P_k}[(Y_i, X_i, W_i) \in \mathcal{A} \cap A_k^c] = \mathbb{P}_{P_k}[(Y_i, X_i, W_i) \in \mathcal{A} \mid X_i \in A_k^c]\mathbb{P}_{P_k}[X_i \in A_k^c]$$

$$
\begin{aligned}
&= \mathbb{P}_Q[(Y_i, X_i, W_i) \in \mathcal{A} \mid X_i \in A_k^c]\mathbb{P}_Q[X_i \in A_k^c] \\
&\to \mathbb{P}_Q[(Y_i, X_i, W_i) \in \mathcal{A}].1
\end{aligned}
$$

by the continuity property of probability measures because $A_k \downarrow \emptyset$ and $A_k^c \uparrow \mathbb{R}$. Therefore, Assumption 1 is satisfied for every $m \in \mu(\mathbb{P})$, and both Theorems 1 and 2 apply. □

The inference procedures suggested by Caetano (2015) rely on nonparametric local polynomial estimation methods. As in the RDD case, these procedures yield tests that are a.s. continuous in the data and confidence intervals of finite expected length. Corollary 3 implies lack of control of size and zero confidence level.

4 Conclusion

When drawing inference on a parameter in econometric models, some papers provide conditions under which tests have trivial power (the first type of impossibility). Others examine when confidence regions can have error probability equal to one (the second type of impossibility). The motivation behind these negative results is that the parameter of interest may be nearly unidentified for some models under the alternative hypothesis. This reasoning requires some notion of distance for the models, in which some authors use the total variation metric (strong convergence) and others rely on the Lévy-Prokhorov metric (weak convergence).

This paper presents necessary and sufficient conditions for tests to have trivial power. This result is of interest because it is needed to establish which notion of model distance is required for power to be bounded by size. Furthermore, the first type of impossibility is stronger than the second type. Dufour (1997) focuses on models in which tests based on bounded confidence regions fail to control size, but they could still have nontrivial power. Take the simultaneous equations model when instrumental variables (IVs) may be arbitrarily weak. Kleibergen (2002) and Moreira (2002, 2003) propose tests that have correct size in models where the second type of impossibility applies. Furthermore, these tests have good power when identification is strong, being efficient under the usual asymptotics. Their power is not trivial, exactly because not every model under the alternative is approximated by models under the null.

We then revisit the impossibility finding of Kamat (2015) in Regression Discontinuity Design models, and study the exogeneity tests of Caetano (2015) based on bunching. One potential solution to the first type of impossibility is to restrict the class of conceivable models to allow power to be larger than size. This is the approach of Armstrong and Kolesar (2015) who restrict regression functions to the class of Lipschitz continuous functions with a bounded Lipschitz constant. They derive minimax optimal confidence intervals of correct size. This optimality result is quite appealing, but applied researchers may find the minimax criterion to be somewhat pessimistic in certain settings. In line with the IV literature, we may want to construct tests with correct size when the first-stage discontinuity may be small, as in Feir, Lemieux, and Marmer (2016) do in the Fuzzy RDD case; but in a restricted class of models, as in Armstrong and Kolesar (2015). As alternative solutions to minimax, we are currently investigating methods that control size based on a given test statistic.

References

Angrist, J., and V. Lavy (1999): "Using Maimonides' Rule to Estimate the Effect of Class Size on Scholastic Achievement," *The Quarterly Journal of Economics*, 114(2), 533–575.

Armstrong, T. B., and M. Kolesar (2015): "Optimal Inference in a Class of Regression Models," *Working Paper, Yale University.*

Bahadur, R. R., and L. J. Savage (1956): "The Nonexistence of Certain Statistical Procedures in Nonparametric Problems," *The Annals of Mathematical Statistics*, 27(4), 1115–1122.

Bertanha, M. (2016): "Regression Discontinuity Design with Many Thresholds," *Available at SSRN, http://dx.doi.org/10.2139/ssrn.2712957.*

Bertanha, M., and G. Imbens (2016): "External Validity in Fuzzy Regression Discontinuity Designs," *CORE Discussion Paper 2016/25.*

Billingsley, P. (2008): *Probability and Measure*. John Wiley & Sons.

Black, S. (1999): "Do Better Schools Matter? Parental Valuation of Elementary Education," *The Quarterly Journal of Economics*, 114(2), 577–599.

Caetano, C. (2015): "A Test of Exogeneity without Instrumental Variables in Models with Bunching," *Econometrica*, 83(4), 1581–1600.

Calonico, S., M. D. Cattaneo, and R. Titiunik (2014): "Robust Nonparametric Confidence Intervals for Regression-Discontinuity Designs," *Econometrica*, 82(6), 2295–2326.

Chamberlain, G. (1987): "Asymptotic Efficiency in Estimation with Conditional Moment Restrictions," *Journal of Econometrics*, 34(3), 305–334.

Dong, Y. (2016): "Jump or Kink? Regression Probability Jump and Kink Design for Treatment Effect Evaluation," *Working Paper, University of California, Irvine*.

Dong, Y., and A. Lewbel (2015): "Identifying the Effect of Changing the Policy Threshold in Regression Discontinuity Models," *The Review of Economics and Statistics*, 97(5), 1081–1092.

Dufour, J.-M. (1997): "Some Impossibility Theorems in Econometrics with Applications to Structural and Dynamic Models," *Econometrica*, 65(6), 1365–1387.

Feir, D., T. Lemieux, and V. Marmer (2016): "Weak Identification in Fuzzy Regression Discontinuity Designs," *Journal of Business & Economic Statistics*, 34(2), 185–196.

Gleser, L. J., and J. T. Hwang (1987): "The Nonexistence of 100(1-α)% Confidence Sets of Finite Expected Diameter in Errors-in-Variables and Related Models," *The Annals of Statistics*, 15(4), 1351–1362.

Hahn, J., P. Todd, and W. Van der Klaauw (2001): "Identification and Estimation of Treatment Effects with a Regression-Discontinuity Design," *Econometrica*, 69(1), 201–209.

Imbens, G., and K. Kalyanaraman (2012): "Optimal Bandwidth Choice for the Regression Discontinuity Estimator," *The Review of Economic Studies*, 79(3), 933–959.

Imbens, G. W., and T. Lemieux (2008): "Regression Discontinuity Designs: A Guide to Practice," *Journal of Econometrics*, 142(2), 615–635.

Jacob, B. A., and L. Lefgren (2004): "Remedial Education and Student Achievement: a Regression-Discontinuity Analysis," *Review of Economics and Statistics*, 86(1), 226–244.

Kamat, V. (2015): "On Nonparametric Inference in the Regression Discontinuity Design," Available at arXiv:1505.06483.

Kleibergen, F. (2002): "Pivotal Statistics for Testing Structural Parameters in Instrumental Variables Regression," *Econometrica*, 70, 1781–1803.

Lee, D. S. (2008): "Randomized Experiments from Non-Random Selection in US House Elections," *Journal of Econometrics*, 142(2), 675–697.

Lehmann, E. L., and J. P. Romano (2005): *Testing Statistical Hypotheses*. Springer Science & Business Media.

Moreira, M. J. (2002): "Tests with Correct Size in the Simultaneous Equations Model," Ph.D. thesis, UC Berkeley.

——— (2003): "A Conditional Likelihood Ratio Test for Structural Models," *Econometrica*, 71, 1027–1048.

Nielsen, H. S., T. Sørensen, and C. Taber (2010): "Estimating the Effect of Student Aid on College Enrollment: Evidence from a Government Grant Policy Reform," *American Economic Journal: Economic Policy*, 2(2), 185–215.

Porter, J. (2003): "Estimation in the Regression Discontinuity Model," *Unpublished Manuscript, University of Wisconsin at Madison*.

Romano, J. P. (2004): "On Non-parametric Testing, the Uniform Behaviour of the t-test, and Related Problems," *Scandinavian Journal of Statistics*, 31(4), 567–584.

Saez, E. (2010): "Do Taxpayers Bunch at Kink Points?," *American Economic Journal: Economic Policy*, 2(3), 180–212.

Schmieder, J. F., T. von Wachter, and S. Bender (2012): "The Effects of Extended Unemployment Insurance Over the Business Cycle: Evidence from Regression Discontinuity Estimates Over 20 Years," *The Quarterly Journal of Economics*, 127(2), 701–752.

Simonsen, M., L. Skipper, and N. Skipper (2016): "Price Sensitivity of Demand for Prescription Drugs: Exploiting a Regression Kink Design," *Journal of Applied Econometrics*, 31(2), 320–337.

Van der Vaart, A. W. (2000): *Asymptotic Statistics*, vol. 3. Cambridge University Press.